

11/28/00  
JC952 U.S. PTO

11-30-00

A notice

<b>UTILITY PATENT APPLICATION TRANSMITTAL</b> <b>(NO FEE)</b> (New Nonprovisional Applications Under 37 CFR § 1.53(b))	Attorney Docket No. 60002-0503
--	-----------------------------------

**TO THE COMMISSIONER FOR PATENTS:**

Transmitted herewith is the patent application of ( ) application identifier or (X) first named inventor, Shai Mohaban, et al., entitled METHOD AND APPARATUS FOR MONITORING AND PROCESSING VOICE OVER INTERNET PROTOCOL PACKETS, for a(n):

- (X) Original Patent Application.
- ( ) Continuing Application (prior application not abandoned):
- ( ) Continuation ( ) Divisional ( ) Continuation-in-part (CIP)  
of prior application No: \_\_\_\_\_ Filed on: \_\_\_\_\_
  - ( ) A statement claiming priority under 35 USC § 120 has been added to the specification.

Enclosed are:

- (X) Patent Application including cover sheet 53 Total Pages; (X) Formal Drawing(s) 10 Sheets ;
- (X) Oath or Declaration: 2 Pages
  - ( ) A Newly Executed Combined Declaration and Power of Attorney:
    - ( ) Signed. (X) Unsigned. ( ) Partially Signed.
  - ( ) A Copy from a Prior Application for Continuation/Divisional (37 CFR § 1.63(d)).
    - ( ) Incorporation by Reference. The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied, is considered as being part of the disclosure of the accompanying application and is hereby incorporated herein by reference.
    - ( ) Signed Statement Deleting Inventor(s) Named in the Prior Application. (37 CFR § 163(d)(2)).
  - ( ) Power of Attorney. (X) Return Receipt Postcard.
  - ( ) Associate Power of Attorney. ( ) A Check in the amount of \$\_\_\_\_\_ for the Filing Fee.
  - ( ) Preliminary Amendment. ( ) Information Disclosure Statement and Form PTO-1449.
  - ( ) A Duplicate Copy of this Form for Processing Fee Against Deposit Account.
  - ( ) A Certified Copy of Priority Documents (if foreign priority is claimed).
  - ( ) Applicant claims small entity status. See 37 CFR 1.27.
  - ( ) Statement(s) of Status as a Small Entity Filed in Prior Application, Status Still Proper and Desired.
  - ( ) Other: \_\_\_\_\_

JC690 U.S. PTO  
09/724943  
09/28/00

CLAIMS AS FILED				
FOR	NO. FILED	NO. EXTRA	RATE	FEE
Total Claims	64	44	\$18.00	\$ 792.00
Independent Claims	11	8	\$80.00	\$ 640.00
Multiple Dependent Claims (if applicable)				\$ 0.00
Basic Filing Fee				\$ 710.00
Total Filing Fee				\$2142.00

**NO FEES WILL BE PAID AT THIS TIME.**

Charge \$\_\_\_\_\_ to Deposit Account \_\_\_\_\_ pursuant to 37 CFR § 1.25. At any time during the pendency of this application, please charge any fees required or credit any overpayment to this Deposit Account.

Respectfully submitted,

By: Carina M. Tan  
Carina M. Tan, Reg. No.45,769

Date: November 28, 2000

Correspondence Address:

Hickman Palermo Truong & Becker, LLP  
1600 Willow Street  
San Jose, California 95125-5106  
Telephone: (408) 414-1080  
Facsimile: (408) 414-1076

I hereby certify that this is being deposited with the U.S. Postal Service "Express Mail Post Office to Addressee" service under 37 CFR § 1.10 on the date indicated below and is addressed to:

Commissioner for Patents  
Box Patent Application  
Washington, D.C. 20231

By: Casey Moore

Typed Name: Casey Moore

Express Mail Label No.: EL652871106us

Date of Deposit: November 28, 2000

60002-0503

Patent

UNITED STATES PATENT APPLICATION  
FOR  
METHOD AND APPARATUS FOR MONITORING AND PROCESSING VOICE OVER INTERNET  
PROTOCOL PACKETS

INVENTORS:

SHAI MOHABAN  
ITZHAK PARNAFES  
OPHER KAHANE

PREPARED BY:

HICKMAN PALERMO TRUONG & BECKER LLP  
1600 WILLOW STREET  
SAN JOSE, CA 95125-5106  
(408) 414-1080

EXPRESS MAIL CERTIFICATE OF MAILING

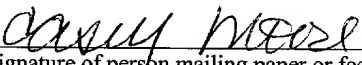
"Express Mail" mailing label number: EL652871106US

Date of Deposit: November 28, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

CASEY MOORE

(Typed or printed name of person mailing paper or fee)



(Signature of person mailing paper or fee)

Variable	Mean	SD	Min	Max
Age	34.2	10.5	20	55
Gender	0.5	0.5	0	1
Marital status	0.6	0.5	0	1
Education	12.5	1.5	9	16
Income	1.2	0.8	0.5	2.5
Health status	0.7	0.4	0	1
Stress level	2.5	1.2	1	4
Life satisfaction	3.5	1.0	1	5
Work satisfaction	3.0	1.0	1	5
Family satisfaction	3.5	1.0	1	5
Community satisfaction	3.0	1.0	1	5
Overall satisfaction	3.2	1.0	1	5

## RELATED APPLICATION

This application claims domestic priority from prior U.S. Provisional application Ser. No. 60/226,207, filed August 18, 2000, entitled “ Method and Apparatus for VoIP Traffic Processing,” and naming as inventors S. Mohaban et al., the entire disclosure of which is hereby incorporated by reference as if fully set forth herein.

## FIELD OF THE INVENTION

The present invention relates generally to data networks that carry voice traffic. The invention relates more specifically to a method and apparatus for monitoring and processing information traveling over Internet Protocol IP networks including voice over IP, video over IP, and streaming media.

## BACKGROUND OF THE INVENTION

Packet-switched data networks now carry a high volume of messages (“ traffic”) pertaining to specialized services such as digitized voice, music, video, and streaming media. There is particular technical interest in improving the capabilities of the global, packet-switched family of internetworks known as the Internet for carrying voice conversations, as an alternative to the traditional circuit-switched public telephone network.

Advancements in voice over Internet Protocol services (VoIP) have lead to development of numerous technical recommendations and protocols. Significantly, VoIP is no longer a technical novelty, but a real business for a growing number of for-profit organizations that sell and service VoIP connectivity (“service providers”). An overview of these developments is provided in U. Black, “Voice Over IP” (Prentice-Hall, 2000).

VoIP service providers now face challenges similar to those experienced by electronic commerce (“eCommerce”) service providers at the onset of the Internet explosion. The typical eCommerce service provider network comprised routers, switches,





## SUMMARY OF THE INVENTION

A processor architecture for processing data packets representing voice over Internet Protocol (VoIP) calls in a packet-switched network is disclosed. According to an embodiment, a VoIP processor executes a voice packet processing operating system that is configured to monitor or manipulate the packets at an IP layer, media layer and signaling layer of the call. The VoIP processor includes a plurality of independently callable primitive software functions that carry out low-level VoIP packet processing functions. The VoIP processor executes one or more application programs that selectively call one or more of the primitive software functions and are independent of any underlying protocols of the existing network, thereby isolating the application programs from low-level processing details. Further, techniques are described for modifying characteristics of VoIP traffic for the purpose of monitoring and directing the VoIP traffic through a network. The techniques include extracting information associated with the VoIP traffic and using the information for the purpose of controlling access, for fraud detection, for billing, for enforcing policy decisions, for protection against denial of service attacks, for lawful interception, for service selection, and other applications.

In one aspect, the invention provides a method of processing data packets representing voice over Internet Protocol (VoIP) calls in a packet-switched network. One or more VoIP processors are provided in the network. Each of the VoIP processors (a) executes a voice packet processing operating system that is configured to monitor or manipulate the packets at a transport layer, media layer and signaling layer of the call, and (b) includes a plurality of independently callable primitive software functions that carry out low-level VoIP packet processing functions. One or more application programs that provide one or more call processing functions by selectively calling one or more of the primitive software functions and are independent of any underlying protocols of the existing network may be executed. In one embodiment, the one or more applications may be executed by the VoIP

processor. In other embodiments, the one or more applications may be executed by other network devices. In operation, the VoIP processor detects VoIP packets that pass through the one or more VoIP processors and identifying one or more values of fields in the packets. The VoIP processor creates and stores call state information associated with each call that is  
5 represented by the one or more VoIP packets. The processor also modifies one or more of the VoIP packets at either the transport layer, the media layer, or the signaling layer as required to carry out one or more call processing functions of the one or more application programs.

Other aspects and features will become apparent from the following description and  
10 the appended claims.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

5           FIG. 1 is a block diagram that illustrates an example location of a VoIP processor within a service provider point of presence;

          FIG. 2 is a block diagram that illustrates an example location of a VoIP processor within an Access POP;

          FIG. 3 is a block diagram showing hardware elements of a VoIP processor in one  
10   example embodiment;

          FIG. 4A is a block diagram that illustrates a software architecture of a VoIP processor according to one embodiment;

          FIG. 4B is a block diagram that illustrates some example VoIP processor primitives;

          FIG. 4C is a block diagram that illustrates some example VoIP application programs ;

15           FIG. 5 is a block diagram that illustrates a virtual path formed by VoIP processors in a network under control of application-level call routing application;

          FIG. 6 is a block diagram that illustrates a network that is overlaid on a general network by using VoIP processors in traffic engineering domains;

          FIG. 7 is a block diagram of an example point of presence in which a VoIP processor  
20   acts as a load balancer for a plurality of gateways; and

          FIG. 8 depicts a computer upon which embodiments of the invention may be implemented.



[illegible]

5

## 15

15

20

25

interface and at least one "outbound" interface, such that network traffic flows bi-directionally between inbound interfaces and outbound interfaces. The physical network interfaces may be any of a plurality of interfaces, e.g., Fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, Packet Over Sonet (POS) interfaces, etc.

5           In this embodiment, the VoIP processor can track VoIP calls as packet flows associated with the calls move through the network. The VoIP processor is able to perform real-time packet manipulation and measurements at a plurality of logical layers. In one embodiment, the VoIP processor can manipulate and measure packet characteristics associated with the IP layer, media layer, and call signaling layer.

10           In another embodiment, a VoIP processor includes a plurality of packet-processing software elements that may be incorporated into application programs that carry out a variety of useful functions. Such protocol tracking and packet processing "primitives" can achieve certain functional goals as described below.

15           In one specific embodiment, a VoIP processor is communicatively coupled in a service provider network at a point of physical or logical discontinuity. For example, a VoIP processor is installed at the exit point from a Point Of Presence (POP) to the Wide Area Network (WAN). Alternatively, the VoIP processor is communicatively coupled between two adjacent networks. In still another alternative, the VoIP processor is communicatively coupled at a strategic peering point within the network.

20           FIG. 1 is a block diagram that illustrates an example location of a VoIP processor within a service provider point of presence.

25           In the example of FIG. 1, VoIP point of presence 120 is communicatively coupled to an Internet Protocol wide area network ("IP WAN") 132 and to one or more end user devices, such as a facsimile machine 112, or a telephone 110 of the type conventionally used with the public switched telephone network ("PSTN phone"). Point of presence 120 comprises a VoIP processor 126 that is communicatively coupled to IP WAN 132 through router 130.

VoIP processor 126 is also communicatively coupled to one or more Gateways, two of which are illustrated in FIG. 1 as examples, namely, Gateways 124a and 124b. Gateways 124a and 124b are communicatively coupled to a PSTN switch 122. PSTN switch 122 is communicatively coupled to facsimile machine 112 and PSTN phone 110. Fax machine 112 and phone 110 are illustrated for example purposes, and in a practical embodiment there may be zero to any number of such devices. In addition, there may be any number of gateways and switches.

A VoIP processor may also be used in an Access point of presence of an Internet Service Provider ("ISP") to support endpoints that are communicatively coupled to the ISP through various connection types, such as dial-up, DSL or cable. FIG. 2 is a block diagram that illustrates an example location of a VoIP processor within an Access POP.

In FIG. 2, Access POP 220 is communicatively coupled to an IP WAN 260. Access POP 220 comprises a VoIP processor 234 that is communicatively coupled to IP WAN 260 through router 250. VoIP processor 234 is also communicatively coupled to a digital subscriber line access multiplexer (DSLAM) 232, a dial-up concentrator 228 and a cable modem termination system (CMTS) 224. DSLAM 232 is communicatively coupled to a message transfer agent (MTA) 230, which is in turn communicatively coupled to consumer devices, IP phone 210 and workstation 208. Dial-up concentrator 228 is communicatively coupled to modem 226, which is communicatively coupled to workstation 206. CMTS 224 is connector MTA 222, which is communicatively coupled to PSTN phone 204 and workstation 202.

## 2. HARDWARE ARCHITECTURE

FIG. 3 is a block diagram showing hardware elements of a VoIP processor in one example embodiment.

In FIG. 3, VoIP processor 300 generally comprises a switching circuitry 302, physical interfaces 304, classification engines 306, classification tables 307, network

processors 308 and a host central processing unit (CPU) 310. Each of the foregoing elements may be implemented in hardware, firmware, software, or a combination thereof.

VoIP packets or cells enter and leave the VoIP processor through physical interfaces 304. An example of a physical interface is an Ethernet transceiver or ATM transceiver.

5 There may be any number of such physical interfaces in a VoIP processor.

Upon entering the VoIP processor, VoIP packets pass through switching circuitry 302, which hands the VoIP packets to other components in the system or possibly back to the network through another physical interface. Switching circuitry 302 may comprise a switch matrix or fabric of the type conventionally found in a router, bridge or switch.

10 Classification engines 306 function to classify the VoIP traffic that enters the VoIP processor by extracting various fields from the VoIP packets. In an embodiment, VoIP processor has one or more pre-processed classification tables 307. When a packet arrives, classification engines 306 examine values of packet header fields such as the protocol field, the source and destination IP address fields, and the source and destination port fields. A  
15 table look-up is made for the extracted fields in the classification tables 307, which yields a tag value for each packet. The VoIP packets are tagged accordingly. Switching circuitry 302 also may use the resulting tags to forward the VoIP packets to other components in the system. In one embodiment, classification engines 306 may also examine the Type of Service field and DSCP value. Relevant packets that have been classified into a pre-  
20 identified protocol such as Q.931, H.245 or SIP also may be parsed to obtain other relevant parameter values. Such parameter values may include lower level protocol port values, such as the port numbers that are negotiated in the Q.931 protocol for use in an associated H.245 stream. Other parameter values include the caller identifier, called number, type of codec, etc. Such deeper protocol parsing can be done either by the host processor or a network  
25 processor.

One or more dedicated network processors 308 execute various packet manipulations and implement primitive software functions of the type described further herein.

Host CPU 310 executes general software based functions such as deep protocol parsing, hardware management, protocols implementation, etc. Host CPU 310 also acts as an overall supervisor of the other elements of the VoIP processor 300. In certain embodiments, a VOIP processor may be multiple host CPUs. Multiple host processors may be used for load sharing or for fulfilling system redundancy objectives.

In one embodiment, the VoIP processor 300 is implemented using a line card architecture in which the physical interfaces, classification engines and network processors are combined on a circuit board called a line card. One or more such line cards are connected to a backplane. Communications among hardware elements on the backplane may be controlled by a bus circuitry or switching circuitry of any conventional type.

VoIP processor 300 also includes one or more non-volatile mass storage devices, one or more main memory devices, and one or more input/output ports of the type provided in conventional computer servers. VoIP processor 300 may include or interface to other hardware and software systems such as database servers, etc.

In one specific embodiment, the VoIP processor is configured to provide high availability by means of two or more redundant power supplies, duplicate fans, duplicate host CPUs, etc. In another embodiment, each hardware element of the VoIP processor has circuitry that enables removal and replacement of each such element while power remains applied and the processor is running ("hot swapping").

### 3. SOFTWARE ARCHITECTURE

#### 3.1 GENERALLY

In one embodiment, host CPU 310 executes two logically separate software systems: a voice packet processing operating system and one or more VoIP application programs. In this configuration, in operation, VoIP processor 300 may be viewed as a vertical, voice-specific, Policy Enforcement Point (PEP) network device. Specifically, VoIP processor 300 is interposed in a VoIP network and receives and analyzes all packets that go through the network. VoIP processor 300 parses all packets that represent VoIP traffic, and tracks the

state of each call. In certain embodiments, the one or more application programs may be executed on one or more entities that are separate from host CPU 310. For example, the one or more application programs may be executed at one or more external Policy Decision Points (PDP) network devices. VoIP processor 300 creates and stores data based on the  
5 parsed packets, at an abstraction layer logically above the underlying VoIP protocols, such as MGCP, the H.323 recommendations published by the International Telecommunications Union (ITU), etc. In this way, the VoIP application programs may be written independently of the VoIP protocol used. The VoIP processor parses and acts on packets at the IP layer, the media layer and the signaling layer.

10 FIG. 4A is a block diagram that illustrates a software architecture of a VoIP processor according to one embodiment.

In FIG. 4A, VoIP processor 300 executes one or more application programs (“VoIP applications”) 404, and a voice packet processing operating system (“OS”) 406. VoIP  
15 applications 404 comprise application programs that carry out functions useful in voice call processing. For example, VoIP applications 404 may comprise Real-Time Protocol (RTP) aggregation application 414, billing verification application 415, VoIP Denial of Service (DoS) protection 417, VoIP load balancing application 420, etc. Functions and operations of such applications are described further herein in the section entitled “Application Examples.”

Voice packet processing operating system 406 comprises a signaling layer 408 that  
20 performs call processing, a media layer 410 that performs media processing and a transport layer 412 that performs IP processing. In one specific embodiment, voice packet processing operating system 406 includes a plurality of primitive function software elements 413 (“primitives”) that carry out low-level packet processing functions. Each of the primitives may be implemented as a subroutine, programmatic object, or other element that layer 408,  
25 410, 412 or an application program can invoke to carry out a desired low-level function. Specific functions and operations of such primitives are described further herein in the section entitled “Primitive Function Software Elements.” In the case where VOIP

In this or another embodiment, the VoIP processor may also be controlled using the Simple Network Management Protocol (SNMP protocol) using SNMP agent 422, or through a command line interface using CLI interface 424, which may be, for example, a serial port, Telnet agent, etc.

At each layer among signaling layer 408, media layer 410, and transport layer 412, the VoIP processor 300 executes one or more basic primitives that act on the actual VoIP packet streams and are performed in real time without affecting the speed of transport of packets through the network of which the VoIP processor is a part. With this architecture, VoIP processor 300 is specifically designed to expose programmable interfaces and allow external logic to control the platform primitives.

FIG. 4B is a block diagram that illustrates some example VoIP processor primitives. In one embodiment, primitive function software elements 413 comprise IP layer primitives 430, media layer primitives 460, and signaling layer primitives 480, as described further in the following sections.

### 5                    3.2.1 IP LAYER PRIMITIVES

IP layer primitives 430 may include a packet dropping routine 432, packet duplication routine 434, packet marking routine 436, Multi-protocol Label Switching (MPLS) label manipulation routine 438, traffic policing routine 440, packet scheduling routine 442, packet re-routing routine 444, tunneling routine 446, encryption routine 448, packet injection routine 10 450, and packet compression routine 452.

Packet dropping routine 432 is called to drop packets based on various parameters in the IP header. This primitive is useful for preventing certain media and signaling streams from entering the network, and may be used by firewall applications, protection against denial of service attacks application, access control applications, etc.

15            Packet duplication routine 434 is called to duplicate and send packets to a new destination, either by changing the destination IP address or by using various encapsulation protocols. This primitive may be used by a lawful interception application, for example, which may direct the duplicated traffic to a real-time tapping or recording system.

Packet marking routine 436 can modify the Type of Service (ToS) field or the 20 DiffServ CodePoint (DSCP) in the IP header of a packet.

Label manipulation routine 438 can label a packet when the destination IP address of the packet indicates that the packet is entering an MPLS domain. Thus, an application that uses label manipulation routine 438 may cause the VoIP processor to operate as an MPLS gateway. Within an MPLS domain, the label may be modified or a new label may be pushed 25 or popped on top of the label stack. Label manipulation also may be used by traffic engineering applications.





Transcoding routine 462 is called to cause the coder/decoder used by VoIP processor 300 to carry voice or fax media to change in real time (on the fly). For example, transcoding routine 462 may be called when the two endpoints of a call do not support the same codec types.

5           RTP aggregation routine 464 may carry out RTP aggregation, which may also be referred to as RTP multiplexing, or RTP trunking. In RTP aggregation, multiple packets of different RTP flows are aggregated together with one single header, provided the packets share a common sub-route. Thus, when properly invoked, the RTP aggregation routine 464 may save a significant amount of bandwidth.

10           Header compression routine 466 is called to compress an IP header, UDP header, and/or RTP header using various header compression techniques.

Media modification routine 468 is called to cause the VoIP processor to decode the voice stream, modify it and re-encode it. Media modification routine 468 also may be used for applications such as ad insertion, voice announcements, etc.

15           Media reconstruction routine 470 is called to reconstruct dropped or lost packets. When packets are dropped from a VoIP call, the dropped packets have a significant effect on voice quality. Most endpoints are incapable of recovering lost packets. VoIP processor 300, however, may use innovative compressed and uncompressed media reconstruction techniques for recovering some of the lost packets and improve the voice quality.

20           Media duplication routine 472 is called to duplicate one or more media streams within an application, such as a lawful interception application, call logging application, etc.

Media re-routing routine 474 is called to re-route one or more media streams to a new destination or re-routed through a particular network path.

### 3.2.3 SIGNALING LAYER PRIMITIVES

25           Signaling layer primitives 480 may comprise a protocol translation routine 482, call detail record generation routine 484, number translation routine 486, drop call routine 488,







detailed description of an example aggregation algorithm is provided in co-pending application Ser. No. NUMBER, filed DATE, entitled “ Method and Apparatus for Call Aggregation,” by named inventors S. Mohaban et al., attorney docket number 60002-505, the entire contents of which is hereby incorporated by reference as if fully set forth herein.

5 Using RTP aggregation application 414, VoIP processor 300 causes a substantial bandwidth reduction, while improving other aspects of the network behavior, such as reducing the load on routers and improving the behavior of queues.

In one embodiment, RTP aggregation application 414 ensures a prescribed limit on the induced latency. For example, RTP streams are partitioned based on certain criteria including the destination subnet of the RTP stream, and the codec of the RTP stream. When the first packet of a trunk arrives at VoIP processor 300, a timer may be started. The timer is set to expire at an allowed maximum value of latency. The allowed maximum value of latency may be user-selected. As packets continue to arrive at VoIP processor 300, the packets are aggregated either when the number of packets reach a pre-selected number or when the timer expires.

Applications that work on multiple voice sessions simultaneously often get more effective as the number of concurrent calls gets larger. In the case of call aggregation, the larger the number of calls sharing a common network sub-route, the better the bandwidth utilization that can be achieved. By aggregating calls originating from multiple endpoints (gateways and IP phones) the call aggregation application on the VoIP processor improves the ability to aggregate more calls. For example, in case of IP phones, aggregation simply is not applicable on a single phone. However, the calls from several IP phones may be aggregated by the call aggregation application running on a VoIP processor.

### 3.3.1.2 APPLICATION LEVEL CALL ROUTING

of real-time quality sensitive traffic, such as VoIP. Using an application-level call routing application 431, a network operator deploys an overlaid network of VoIP processors. As a result, the network operator can control the routing of the call in an IP network by designing or engineering alternative routes in the network. Such alternative routes can be used when congestion is identified along one of the default IP routes. When congestion is identified, the VoIP processors form a virtual path along which the traffic is redirected, thus circumventing the congested areas in the network.

FIG. 5 is a block diagram that illustrates a virtual path formed by VoIP processors in a network under control of application-level call routing application 431. In FIG. 5, an IP WAN 502 is communicatively coupled to VoIP points of presence 526, 524 by routers 506, 512 and routers 508, 510, respectively. Point of presence 524 includes VoIP processor 520 and VoIP gateway 522. Point of presence 526 includes VoIP processor 514 and VoIP gateway 516. IP WAN 502 includes VoIP processor 504. Each of the VoIP processors 514, 520 is a processor of the type described herein in connection with FIG. 3, FIG. 4A, FIG. 4B, FIG. 4C and each such processor executes a copy of application-level call routing application 431. A virtual path formed by VoIP processors 520, 504, 514 is indicated by the bi-directional arrows 550a, 550b.

In this configuration, the application-level call routing application 431 improves overall call quality by avoiding congested areas. Further, application 431 balances and averages out instantaneous peak loads in the network. In general, Internet data traffic tends to be bursty and unpredictable. Certain links can become congested, which can cause voice packets to be delayed or dropped. As a result, call quality can degrade, or a call can fail to start. In response to such conditions, application 431 can divert network traffic from congested areas of the network to less congested areas of the network. The application 431 improves the predictability of voice service quality by ensuring that the best virtual path is available and placed in use when the network is loaded. The application 431 works in





In one embodiment, using one or more VoIP processors, a network is created or overlaid on top of the general network. VoIP processors are deployed at a domain edge and at peering points. Within each domain, either regular traffic engineering methods are used, such as explicit route establishment using a protocol such as RSVP -TE and CR-LDP, or new routes are created between the VoIP processors within the domain.

FIG. 6 is a block diagram that illustrates a network that is overlaid on a general network by using VoIP processors in traffic-engineered domains, VoIP points of presence and Peering POPs.

In FIG. 6, network 600 comprises Traffic Engineered Domains (e.g., using MPLS) 602, 610, IP networks 606, 614, VoIP points of presence 604, 612 and Peering POPs 608, 618. Traffic Engineered Domain 602 has a VoIP processor 620 located at its logical edge. Traffic Engineered Domain 602 is communicatively coupled to VoIP point of presence 604 and Peering POP 618 by routers 634, 636 and routers 656, 654 respectively. VoIP point of presence 604 includes a VoIP processor 622 that is communicatively coupled to VoIP gateway 624. VoIP point of presence 604 is also communicatively coupled to IP network 606 through routers 636, 638.

IP network 606 is communicatively coupled to Peering POP 608 through router 640, 642. Peering POP 608 includes VoIP processor 626. Peering POP 608 is communicatively coupled to Traffic Engineered MPLS Domain 610 through routers 642, 644. Traffic Engineered MPLS Domain 610 is communicatively coupled to VoIP point of presence 612 through routers 646, 648. VoIP point of presence 612 includes VoIP processor 628 communicatively coupled to VoIP gateway 630. VoIP point of presence 612 is communicatively coupled to IP network 614 through routers 648, 650. IP network 614 is communicatively coupled to Peering POP 618 through routers 652, 654. Peering POP 618 includes VoIP processor 632. Arrows 660, 670 indicate routes between VoIP processors 632, 620, 622.



network path for a specific flow of traffic. However, a problem with RSVP is its inability to scale (handle an unlimited increase in the number of reservation requests) because each router along a network path is required to maintain state information for each microflow that passes through the router. This requirement can mean that certain routers positioned in the backbone of the network must store millions of state information, which is undesirable given the limited storage and processing capabilities of typical routers. However, a recent proposal by the IETF would amend the RSVP specification to allow for aggregation of multiple smaller reservations into one larger reservation in a hierarchical way.

In one embodiment, a VoIP processor 300 carries out such aggregated reservation, and is an ideal location to serve as an aggregation point, especially when operating in conjunction with the application-level call routing application 431. Such larger reservations are established across the network between pairs of VoIP processors according to policies relying on various parameters such as the number of current calls, statistics on call patterns, etc. Such aggregated reservations may be dynamically modified according to network conditions, time of day, expected traffic patterns, etc. The VoIP processor 300 also may also serve as an admission control point in order to decide which calls get to use which routes and reservations.

### 3.3.1.5 LOAD BALANCING AMONG VOICE GATEWAYS

VoIP points of presence often have a multitude of VoIP gateways, and new gateways are added as more capacity is required. Managing such large farms of VoIP gateways becomes complex and does not allow the ISP to fully exploit the available resources.

FIG. 7 is a block diagram of an example point of presence in which a VoIP processor acts as a load balancer for a plurality of gateways. In one embodiment, VoIP point of presence 702 includes any number of VoIP gateways 704A, 704B, 704N. A VoIP processor 300 is communicatively coupled to the VoIP gateways 704A, 704B, 704N and is logically located between the gateways and a router 706. In this configuration, by deploying a VoIP

processor in front of a gateway farm, an ISP can configure the entire farm to be represented to the rest of the network by a single IP address, namely the IP address of the VoIP processor. Further, VoIP processor 300 balances the load between the VoIP gateways. The VoIP processor can support any of a plurality of conventional load-balancing algorithms of the type used to load balance between Web servers in Web server farms.

In this configuration, VoIP processor 300 also can decide, in real time and based on dynamic measurements and provisioned information, how to route an incoming call and make sure it is handled by the most appropriate gateway in the farm according to various policies set by the network administrator.

### 3.3.2 REVENUE PROTECTION APPLICATIONS

VoIP processor 300 can execute any of a plurality of revenue protection applications , including a billing & interconnect verification application 439, access control and fraud detection application 441, RTP reconstruction application 443, and differentiated services application 447.

#### 3.3.2.1 BILLING & INTERCONNECT VERIFICATION

As described above, in one embodiment, VoIP processor 300 parses each packet to identify values of fields in VoIP protocols at the media layer and the signaling layer. Accordingly, VoIP processor 300 can collect accurate details about each voice call that is traversing a network in which the VoIP processor resides. Such details may include the call length, number of packets transmitted, service granted to such packets (e.g., a DiffServ CodePoint or RSVP reservation), quality (e.g. jitter, dropped packets), etc.

In one embodiment, VoIP processor 300 creates and stores such information in a database. Under control of a billing & interconnect verification application 424, VoIP processor issues a detailed, accurate and vendor independent Call Details Record (CDR or IPDR) containing such information.

Unlike a Gatekeeper, VoIP processor 300 monitors not only signaling layer information but also media layer information, and therefore VoIP processor 300 can





ear is sensitive to such loss and even a loss of a fraction of the packets can lead to an unacceptable quality, depending on the actual codec that is in use. However, most of a lost RTP packet can be recovered and reconstructed based on various extrapolation techniques. Such techniques can work on either compressed voice media packets, or on uncompressed analog voice information. However, most endpoints do not have the CPU power or sufficient programmed intelligence to be able to recover such lost packets.

In one embodiment, a VoIP processor is positioned logically in front of a call endpoint and executes an RTP reconstruction application 443. In this configuration, the VoIP processor serves as an RTP reconstructor for such endpoints. For example, the VoIP processor can sense lost packets, and in response, use various techniques and algorithms for reconstructing a packet that is similar to the original missing packet. The artificially generated packet is then sent to the endpoint. As a result, voice quality is improved.

#### 3.3.2.5 DIFFERENTIATED SERVICES

In another embodiment, VoIP processor 300 executes a differentiated services application 447. In this configuration, the VoIP processor serves as a Voice Policy Enforcement Point, and can carry out admission control, policing, and ensuring the level of service provided to each call. The VoIP processor can actively set the priority of each call, according to industry standards such as DiffServ and MPLS. It can also accept or deny resource reservation requests based on dynamic or static policies and even serve as an RSVP proxy. VoIP processor 300 also can police and shape packets that pass through it to conform with predefined bandwidth constraints and minimum guarantees.

#### 3.3.3 NEW REVENUE GENERATION APPLICATIONS

VoIP processor 300 can execute any of a plurality of new revenue generation applications, including a call privacy application 445, lawful interception application 449, and service selection application 451.

##### 3.3.3.1 CALL PRIVACY





cooperation with VoIP processor 300, service selection application 451 enables a user to choose between a multitude of services and service providers. For example, a user may dial a specific access code to select a service or VoIP processor 300 may carry out an interactive voice response (IVR) dialogue with a user to select a service. Alternatively, a user may pre-select a service. In response to the user selection information, VoIP processor 300 can route the calls to the appropriate carrier. The access service provider can route the calls through a private high quality network, or through the open Internet based on how much the user is willing to pay. Also, the access service provider can work with multiple carriers offering different service levels and pass these services on to the end user.

#### 4. HARDWARE OVERVIEW

FIG. 8 is a block diagram that illustrates a computer system 800, elements of which may be used in an embodiment of VoIP processor 300.

Computer system 800 includes a bus 802 or other communication mechanism for communicating information, and a processor 804 coupled with bus 802 for processing information. Computer system 800 also includes a main memory 806, such as a random access memory (RAM), flash memory, or other dynamic storage device, coupled to bus 802 for storing information and instructions to be executed by processor 804. Main memory 806 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 804. Computer system 800 further includes a read only memory (ROM) 808 or other static storage device coupled to bus 802 for storing static information and instructions for processor 804. A storage device 810, such as a magnetic disk, flash memory or optical disk, is provided and coupled to bus 802 for storing information and instructions.

A communication interface 818 may be coupled to bus 802 for communicating information and command selections to processor 804. Interface 818 is a conventional serial interface such as an RS-232 or RS-422 interface. An external terminal 812 or other computer system connects to the computer system 800 and provides commands to it using



take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 810. Volatile media includes dynamic memory, such as main memory 806. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 802. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications.

Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 804 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 800 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector coupled to bus 802 can receive the data carried in the infrared signal and place the data on bus 802. Bus 802 carries the data to main memory 806, from which processor 804 retrieves and executes the instructions. The instructions received by main memory 806 may optionally be stored on storage device 810 either before or after execution by processor 804.

Communication interface 818 also provides a two-way data communication coupling to a network link 820 that is connected to a local network 822. For example, communication interface 818 may be an integrated services digital network (ISDN) card or a



## 5. CONCLUSION

A VoIP processor, primitive function software elements, and associated application programs are disclosed herein. In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that  
5 various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

Beneficially, the VoIP processor provides abstraction of underlying protocols. Specifically, the VoIP processor can hide from the applications the details that are protocol  
10 specific, thus creating an abstraction layer. The abstraction allows applications to focus on the specific logic that they implement rather than on non-important details. Such abstraction can scale to new protocols as they evolve over time.

Further, the VoIP processor provides a flexible voice PEP with a set of voice PDP applications controlling the primitives. This architecture allows for very efficient primitive  
15 implementation and enables external applications containing the actual logic to be developed and deployed rapidly on top of it. Another benefit of the VoIP processor is its focus on voice traffic rather than on general data traffic. The primitives are structured specifically for voice applications.

---

## CLAIMS

What is claimed is:

- 1 1. A method of processing data packets representing voice over Internet Protocol  
2 (VoIP) calls in a packet-switched network, the method comprising the computer-  
3 implemented steps of:  
4 providing in the packet-switched network one or more VoIP processors of a plurality  
5 of VoIP traffic processing devices wherein each of the VoIP processors (a)  
6 executes a voice packet processing operating system that is configured to  
7 monitor or manipulate voice packets at a transport layer, a media layer and a  
8 signaling layer of the call, and (b) includes a plurality of independently  
9 callable primitive software functions that carry out low-level VoIP packet  
10 processing functions;  
11 detecting VoIP packets that pass through the one or more VoIP processors and  
12 identifying one or more values of fields in the packets;  
13 creating and storing call state information associated with each call that is represented  
14 by the one or more VoIP packets; and  
15 modifying one or more of the VoIP packets at either the transport layer, the media  
16 layer, or the signaling layer as required to carry out one or more call  
17 processing functions of one or more application programs.
- 1 2. The method of Claim 1, wherein one or more of the plurality of independently  
2 callable primitive software functions are selectively called by the one or more  
3 application programs that provide the one or more call processing functions and are  
4 independent of any underlying protocols of an existing network, and wherein the one  
5 or more application programs are executed on one or more devices that are separate  
6 from the one or more VoIP processors.
- 1 3. The method of Claim 1, wherein one or more of the plurality of independently  
2 callable primitive software functions are selectively called by the one or more  
3 application programs that provide the one or more call processing functions and are



- 1 10. The method of Claim 1, wherein one of the plurality of independently callable  
2 primitive software functions comprises injecting at the signaling layer new VoIP  
3 packets to augment the one or more VoIP packets.
- 1 11. The method of Claim 1, wherein one of the plurality of independently callable  
2 primitive software functions comprises modifying at the media layer one or more  
3 media characteristics of a plurality of media characteristics associated with the one or  
4 more VoIP packets.
- 1 12. The method of Claim 1, wherein one of the plurality of independently callable  
2 primitive software functions comprises modifying at the signaling layer one or more  
3 signaling characteristics of a plurality of signaling characteristics associated with the  
4 one or more VoIP packets.
- 1 13. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises duplicating the one or more VoIP packets at the  
3 transport layer.
- 1 14. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises marking the one or more VoIP packets at the  
3 transport layer.
- 1 15. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises labeling the one or more VoIP packets at the  
3 transport layer.
- 1 16. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises scheduling for routing the one or more VoIP packets  
3 at the transport layer.



- 1 17. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises re-routing the one or more VoIP packets at the  
3 transport layer.
- 1 18. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises tunneling the one or more VoIP packets at the  
3 transport layer.
- 1 19. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises encrypting the one or more VoIP packets at the  
3 transport layer.
- 1 20. The method of Claim 4, wherein modifying at the transport layer the one or more  
2 VoIP packets further comprises compressing the one or more VoIP packets at the  
3 transport layer.
- 1 21. The method of Claim 11, wherein modifying at the media layer one or more media  
2 characteristics further comprises changing a codec associated with the one or more  
3 VoIP packets.
- 1 22. The method of Claim 11, wherein modifying at the media layer one or more media  
2 characteristics further comprises aggregating the one or more VoIP packets under a  
3 single Real-Time Transport Protocol header when the one or more VoIP packets  
4 share a common sub-route.
- 1 23. The method of Claim 11, wherein modifying at the media layer one or more media  
2 characteristics further comprises compressing headers associated with the one or  
3 more VoIP packets.

1     24.     The method of Claim 11, wherein modifying at the media layer one or more media  
2             characteristics further comprises de-coding a media stream associated with the one or  
3             more VoIP packets.

1     25.     The method of Claim 11, wherein modifying at the media layer one or more media  
2             characteristics further comprises re-coding a media stream associated with the one or  
3             more VoIP packets

1     26.     The method of Claim 11, wherein modifying at the media layer one or more media  
2             characteristics further comprises re-constructing a media stream associated with the  
3             one or more VoIP packets.

1     27.     The method of Claim 11, wherein modifying at the media layer one or more media  
2             characteristics further comprises duplicating a media stream associated with the one  
3             or more VoIP packets.

1     28.     The method of Claim 11, wherein modifying at the media layer one or more media  
2             characteristics further comprises re-routing a media stream associated with the one or  
3             more VoIP packets.

1     29.     The method of Claim 12, wherein modifying at the signaling layer one or more  
2     signaling characteristics further comprises translating one or more signaling  
3     protocols from a plurality of signaling protocols associated with the signaling layer.

1     30.     The method of Claim 12, wherein modifying at the signaling layer one or more  
2     signaling characteristics further comprises generating call detail records and IP detail  
3     records by parsing VoIP protocols at an IP layer, the media layer and the signaling  
4     layer.

1     31.     The method of Claim 12, wherein modifying at the signaling layer one or more  
2     signaling characteristics further comprises translating one or more signaling fields in  
3     a signaling message associated with the one or more VoIP packets.

- 1 32. The method of Claim 12, wherein modifying at the signaling layer one or more  
2 signaling characteristics further comprises performing resource reservation for a call  
3 signal.
- 1 33. The method of Claim 12, wherein modifying at the signaling layer one or more  
2 signaling characteristics further comprises re-routing a call signal.
- 1 34. The method of Claim 12, wherein modifying at the signaling layer one or more  
2 signaling characteristics further comprises re-directing a call signal based on a load  
3 balancing criteria for determining a use of gateways.
- 1 35. The method of Claim 12, wherein modifying at the signaling layer one or more  
2 signaling characteristics further comprises duplicating call signals.
- 1 36. The method of Claim 12, wherein modifying at the signaling layer one or more  
2 signaling characteristics further comprises aggregating call signals.
- 1 37. The method of Claim 36, further comprises maintaining an open signaling  
2 connection between two or more VoIP processors of the plurality of VoIP traffic  
3 processing devices.
- 1 38. The method of Claim 1, wherein the one or more application programs implement  
2 one or more policy decisions associated with a VoIP traffic.
- 1 39. The method of Claim 1, wherein the one or more application programs comprise a  
2 call aggregation application program that aggregates a plurality of packets pertaining  
3 to a same call into one aggregated packet.
- 1 40. The method of Claim 1, wherein the one or more application programs comprise an  
2 application level call routing application and further comprising the steps of:

providing one of the VoIP processors at each of a plurality of alternative routing points in the packet-switched network such that the VoIP processors form a virtual network overlaid upon the packet-switched network;

at one of the VoIP processors, detecting network congestion along a conventional routing path in the packet-switched network;

re-routing a call along a virtual path between two or more of the VoIP processors of the virtual network.

41. The method of Claim 1, wherein the one or more application programs comprise a traffic engineering application further comprising the steps of:

providing one of the VoIP processors at each of a plurality of alternative routing points in the packet-switched network such that the VoIP processors form a virtual network that is overlaid upon the packet-switched network, wherein a first one of the alternative routing points is at an edge of a traffic-engineered domain of the packet-switched network and at least a second of the alternative routing points is at a peering point of presence associated with a service provider of the network;

under control of the traffic engineering application,

creating and storing one or more specified routes from at least the first one of the alternative routing points to at least the second of the alternative routing points, in association with network performance criteria information;

at one of the VoIP processors, detecting that the network performance criteria are satisfied for at least one call that has been routed along a conventional routing path in the packet-switched network;

re-routing the at least one call along one of the specified routes between two or more of the VoIP processors of the virtual network.

42. The method of Claim 1, wherein the one or more application programs comprise a resource reservation application, and further comprising the steps of:

providing one of the VoIP processors at each of a plurality of resource reservation

points in the packet-switched network;

under control of the resource reservation application,

at one of the VoIP processors, detecting that one of a plurality of network

performance criteria are satisfied for at least one call that has been routed

along a conventional routing path in the packet-switched network;





4 forwarding the data packets in encrypted form to a second VoIP processor that is  
5 associated with a terminating endpoint of the call and that is configured to  
6 decrypt the data packets.

1 52. The method of Claim 1, wherein the one or more application programs comprise a  
2 differentiated services application and further comprising the steps of, under control  
3 of the differentiated services application, carrying out one or more operations  
4 selected from among: setting a priority value of all packets associated with a  
5 specified call; accepting a resource reservation request based on one or more  
6 policies; denying the resource reservation request based on one or more policies.

1 53. The method of Claim 1, wherein the one or more application programs comprise a  
2 lawful interception application and further comprising the steps of:  
3 under control of the lawful interception application, duplicating one or more packets  
4 associated with a specified call that is a subject of law enforcement  
5 monitoring to result in creating and storing a plurality of duplicated packets;  
6 forwarding the duplicated packets to a tapping application that is configured to  
7 reconstruct the specified call based on the duplicated packets.

1 54. The method of Claim 1, wherein the one or more application programs comprise a  
2 service selection application and further comprising the steps of:  
3 logically coupling the one or more VoIP processor to a plurality of service providers  
4 that can send and receive calls destined or originating outside the packet-  
5 switched network;  
6 receiving service selection information that identifies one of the plurality of service  
7 providers that has been previously selected by a calling party associated with  
8 a specified call that is passing through the VoIP processor;  
9 forwarding all packets associated with the specified call to one of the plurality of  
10 service providers based on the service selection information.

1 55. A computer-readable medium comprising one or more sequences of instructions for  
2 of processing data packets representing voice over Internet Protocol (VoIP) calls in a  
3 packet-switched network, wherein the packet-switched network includes one or more  
4 VoIP processors wherein each of the VoIP processors (a) executes a voice packet

processing operating system that is configured to monitor or manipulate one or more voice packets at a transport layer, a media layer and a signaling layer of the call, (b) includes a plurality of independently callable primitive software functions that carry out low-level VoIP packet processing functions, and (c) executes one or more application programs that provide one or more call processing functions by selectively calling one or more of the primitive software functions and are independent of any underlying protocols of an existing network, and in which the sequences of instructions, when executed by one or more processors, cause the one or more processors to carry out the steps of:

- detecting VoIP packets that pass through the one or more VoIP processors and identifying one or more values of fields in the packets;
- creating and storing call state information associated with each call that is represented by the one or more VoIP packets;
- modifying one or more of the VoIP packets at either the transport layer, the media layer, or the signaling layer as required to carry out one or more call processing functions of the one or more application programs.

56. An apparatus for processing data packets representing voice over Internet Protocol (VoIP) calls in a packet-switched network, comprising:  
one or more VoIP processors in the packet-switched network, wherein each of the VoIP processors (a) executes a voice packet processing operating system that is configured to monitor or manipulate the packets at a transport layer, a media layer and a signaling layer of the call, (b) includes a plurality of independently callable primitive software functions that carry out low-level VoIP packet processing functions, and (c) executes one or more application programs that provide one or more call processing functions by selectively calling one or more of the primitive software functions and are independent of any underlying protocols of an existing network;  
means for detecting VoIP packets that pass through the one or more VoIP processors and identifying one or more values of fields in the packets;  
means for creating and storing call state information associated with each call that is represented by the one or more VoIP packets;





15 tracking and storing a state information associated with the one or more VoIP  
16 packets; and  
17 executing one or more VoIP applications of a plurality of VoIP applications  
18 that are independent of any underlying protocols of the network.

1 59. An apparatus for processing VoIP traffic on a network, the apparatus comprising:  
2 one or more physical interfaces through which VoIP packets enter and leave;  
3 a switching interface that receives the VoIP packets from the one or more physical  
4 interfaces for distribution to one or more components of a VoIP system;  
5 one or more classification engines coupled to the switching interface for classifying  
6 the VoIP traffic;  
7 one or more processors coupled to the switching interface and to the one or more  
8 classification engines to receive the VoIP packets therefrom;  
9 a memory accessible to the one or more processors; and  
10 one or more sequences of instructions stored in the memory which, when executed  
11 by the one or more processors, cause the one or more processors to carry out  
12 the steps of:  
13 IP processing associated with one or more VoIP packets of a plurality of  
14 VoIP packets at a transport layer;  
15 media processing at a media layer; and  
16 call processing at a signaling layer.

1 60. A method of processing data packets representing video over Internet Protocol data  
2 in a packet-switched network, the method comprising the computer-implemented  
3 steps of:  
4 providing in the packet-switched network one or more video over Internet Protocol  
5 processors wherein each of the video over Internet Protocol processors (a)  
6 executes a video packet processing operating system that is configured to  
7 monitor or manipulate the packets at a transport layer, a media layer and a  
8 signaling layer, (b) includes a plurality of independently callable primitive  
9 software functions that carry out low-level video over Internet Protocol  
10 packet processing functions, and (c) executes one or more application  
11 programs that provide one or more video processing functions by selectively

12 calling one or more of the primitive software functions and are independent  
13 of any underlying protocols of an existing network;  
14 detecting video over Internet Protocol packets that pass through the one or more  
15 video over Internet Protocol processors and identifying one or more values of  
16 fields in the packets;  
17 creating and storing video state information associated with each video display that is  
18 represented by the one or more video over Internet Protocol packets;  
19 modifying one or more of the video over Internet Protocol packets at either the  
20 transport layer, the media layer, or the signaling layer as required to carry out  
21 one or more video processing functions of the one or more application  
22 programs.

1 61. An apparatus for processing data packets representing video over Internet Protocol  
2 data in a packet-switched network, comprising:  
3 one or more video over Internet Protocol data processors in the packet-switched  
4 network, wherein each of the video over Internet Protocol data processors (a)  
5 executes a video packet processing operating system that is configured to  
6 monitor or manipulate the packets at a transport layer, a media layer and a  
7 signaling layer, (b) includes a plurality of independently callable primitive  
8 software functions that carry out low-level video over Internet Protocol data  
9 packet processing functions, and (c) executes one or more application  
10 programs that provide one or more video processing functions by selectively  
11 calling one or more of the primitive software functions and are independent  
12 of any underlying protocols of an existing network;  
13 means for detecting video over Internet Protocol data packets that pass through the  
14 one or more video over Internet Protocol data processors and identifying one  
15 or more values of fields in the packets;  
16 means for creating and storing video state information associated with each video  
17 display that is represented by the one or more video over Internet Protocol  
18 data packets;  
19 means for modifying one or more of the video over Internet Protocol data packets at  
20 either the transport layer, the media layer, or the signaling layer as required to

21 carry out one or more video processing functions of the one or more  
22 application programs.

1 62. An apparatus for processing video over Internet Protocol data traffic, the apparatus  
2 comprising:  
3 means for overlaying one or more video over Internet Protocol data processors of a  
4 plurality of video over Internet Protocol data processors on an existing  
5 network;  
6 means for listening to video over Internet Protocol packets that are passing through  
7 the one or more video over Internet Protocol data traffic processing devices;  
8 means for parsing one or more video over Internet Protocol data packets that are  
9 passing through the one or more video over Internet Protocol data traffic  
10 processing devices;  
11 means for tracking and storing a video state information associated with the one or  
12 more video over Internet Protocol data packets; and  
13 means for providing on the one or more video over Internet Protocol data processors  
14 a video over Internet Protocol data operating system on which can be  
15 executed a plurality of video over Internet Protocol data applications that are  
16 independent of any underlying protocols of the existing network.

1 63. An apparatus for processing video over Internet Protocol data traffic on a network,  
2 the apparatus comprising:  
3 one or more physical interfaces through which video over Internet Protocol data  
4 packets enter and leave;  
5 a switching interface that receives the video over Internet Protocol data packets from  
6 the one or more physical interfaces for distribution to one or more  
7 components of a video over Internet Protocol data system;  
8 one or more classification engines coupled to the switching interface for classifying  
9 the video over Internet Protocol data traffic;  
10 one or more processors coupled to the switching interface and to the one or more  
11 classification engines to receive the video over Internet Protocol data packets  
12 therefrom;

a memory accessible to the one or more processors; and  
one or more sequences of instructions stored in the memory which, when executed by the one or more processors, cause the one or more processors to carry out the steps of:

- listening to the video over Internet Protocol data packets that are passing through the apparatus;
- parsing one or more video over Internet Protocol data packets that are passing through the apparatus;
- tracking and storing a video state information associated with the one or more video over Internet Protocol data packets; and
- executing one or more video over Internet Protocol data applications of a plurality of video over Internet Protocol data applications that are independent of any underlying protocols of the network.

64. An apparatus for processing video over Internet Protocol data traffic on a network, the apparatus comprising:

- one or more physical interfaces through which video over Internet Protocol data packets enter and leave;
- a switching interface that receives the video over Internet Protocol data packets from the one or more physical interfaces for distribution to one or more components of a video over Internet Protocol data system;
- one or more classification engines coupled to the switching interface for classifying the video over Internet Protocol data traffic;
- one or more processors coupled to the switching interface and to the one or more classification engines to receive the video over Internet Protocol data packets therefrom;
- a memory accessible to the one or more processors; and
- one or more sequences of instructions stored in the memory which, when executed by the one or more processors, cause the one or more processors to carry out the steps of:
  - IP processing associated with one or more video over Internet Protocol data packets of a plurality of video over Internet Protocol data packets at a transport layer;



## ABSTRACT OF THE DISCLOSURE

A processor architecture for processing data packets representing voice over Internet Protocol (VoIP) calls in a packet-switched network is disclosed. According to an embodiment, a VoIP processor executes a voice packet processing operating system that is

5 configured to monitor or manipulate the packets at an IP layer, media layer and signaling layer of the call. The VoIP processor includes a plurality of independently callable primitive software functions that carry out low-level VoIP packet processing functions. The VoIP processor executes one or more application programs that selectively call one or more of the primitive software functions and are independent of any underlying protocols of the existing

10 network, thereby isolating the application programs from low-level processing details. Further, techniques are described for modifying characteristics of VoIP traffic for the purpose of monitoring and directing the VoIP traffic through a network. The techniques include extracting information associated with the VoIP traffic and using the information for the purpose of controlling access, for fraud detection, for billing, for enforcing policy

15 decisions, for protection against denial of service attacks, for lawful interception, for service selection, and other applications.

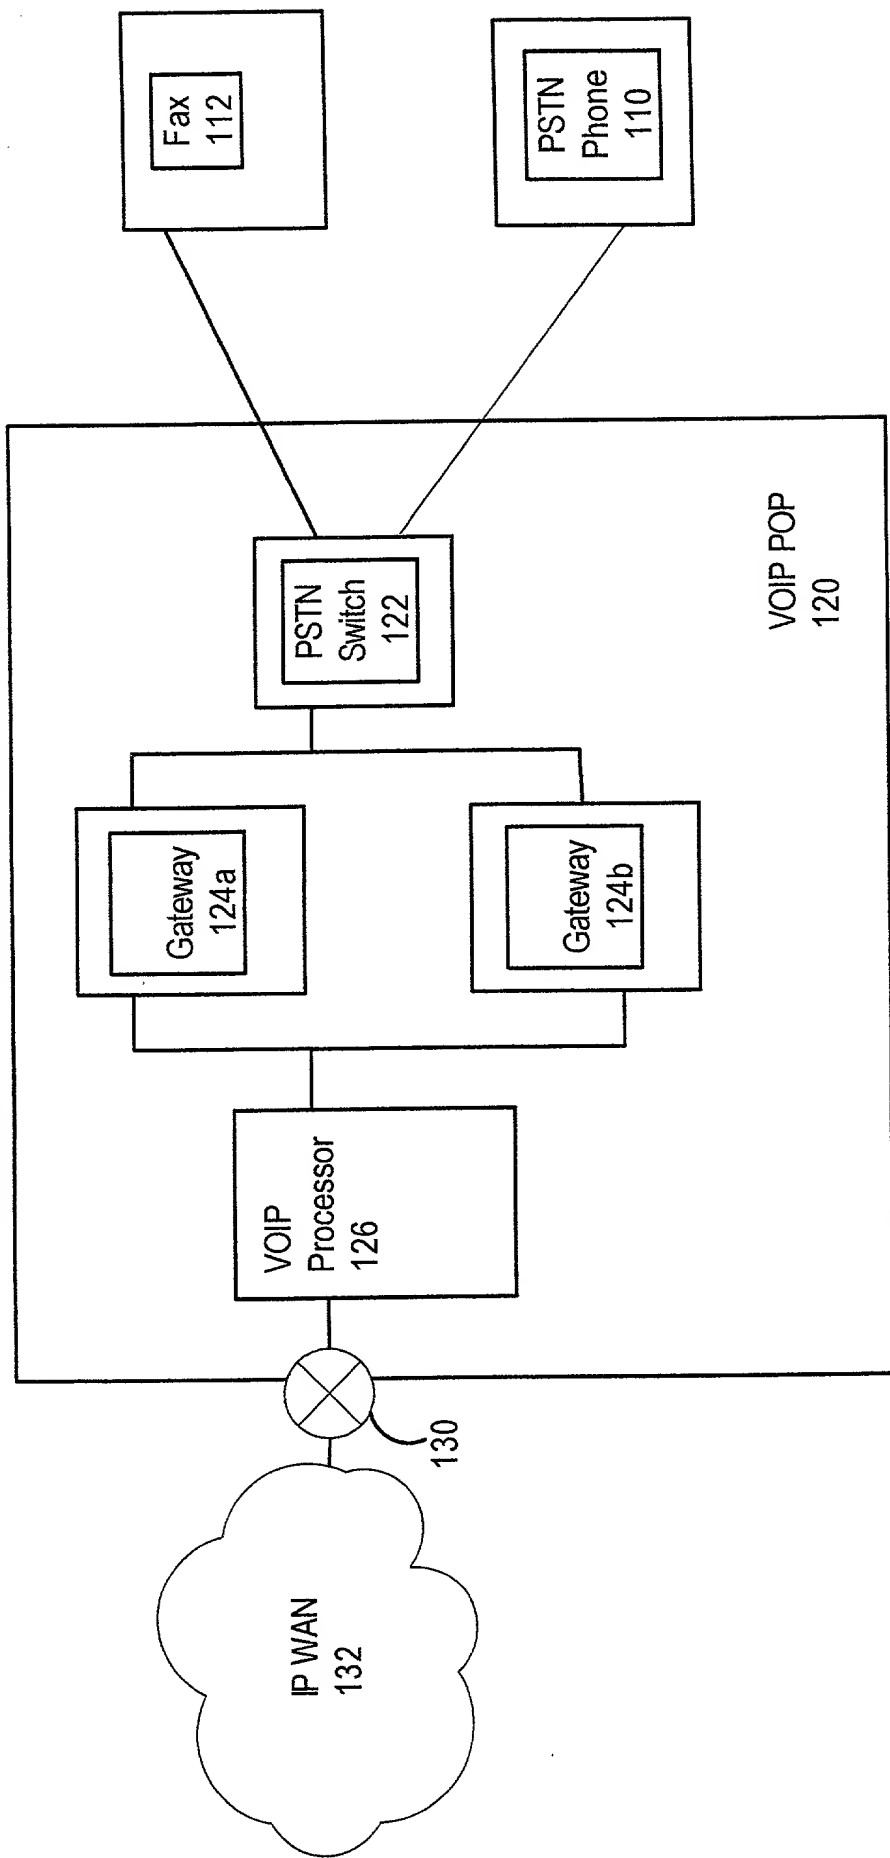


FIG. 1



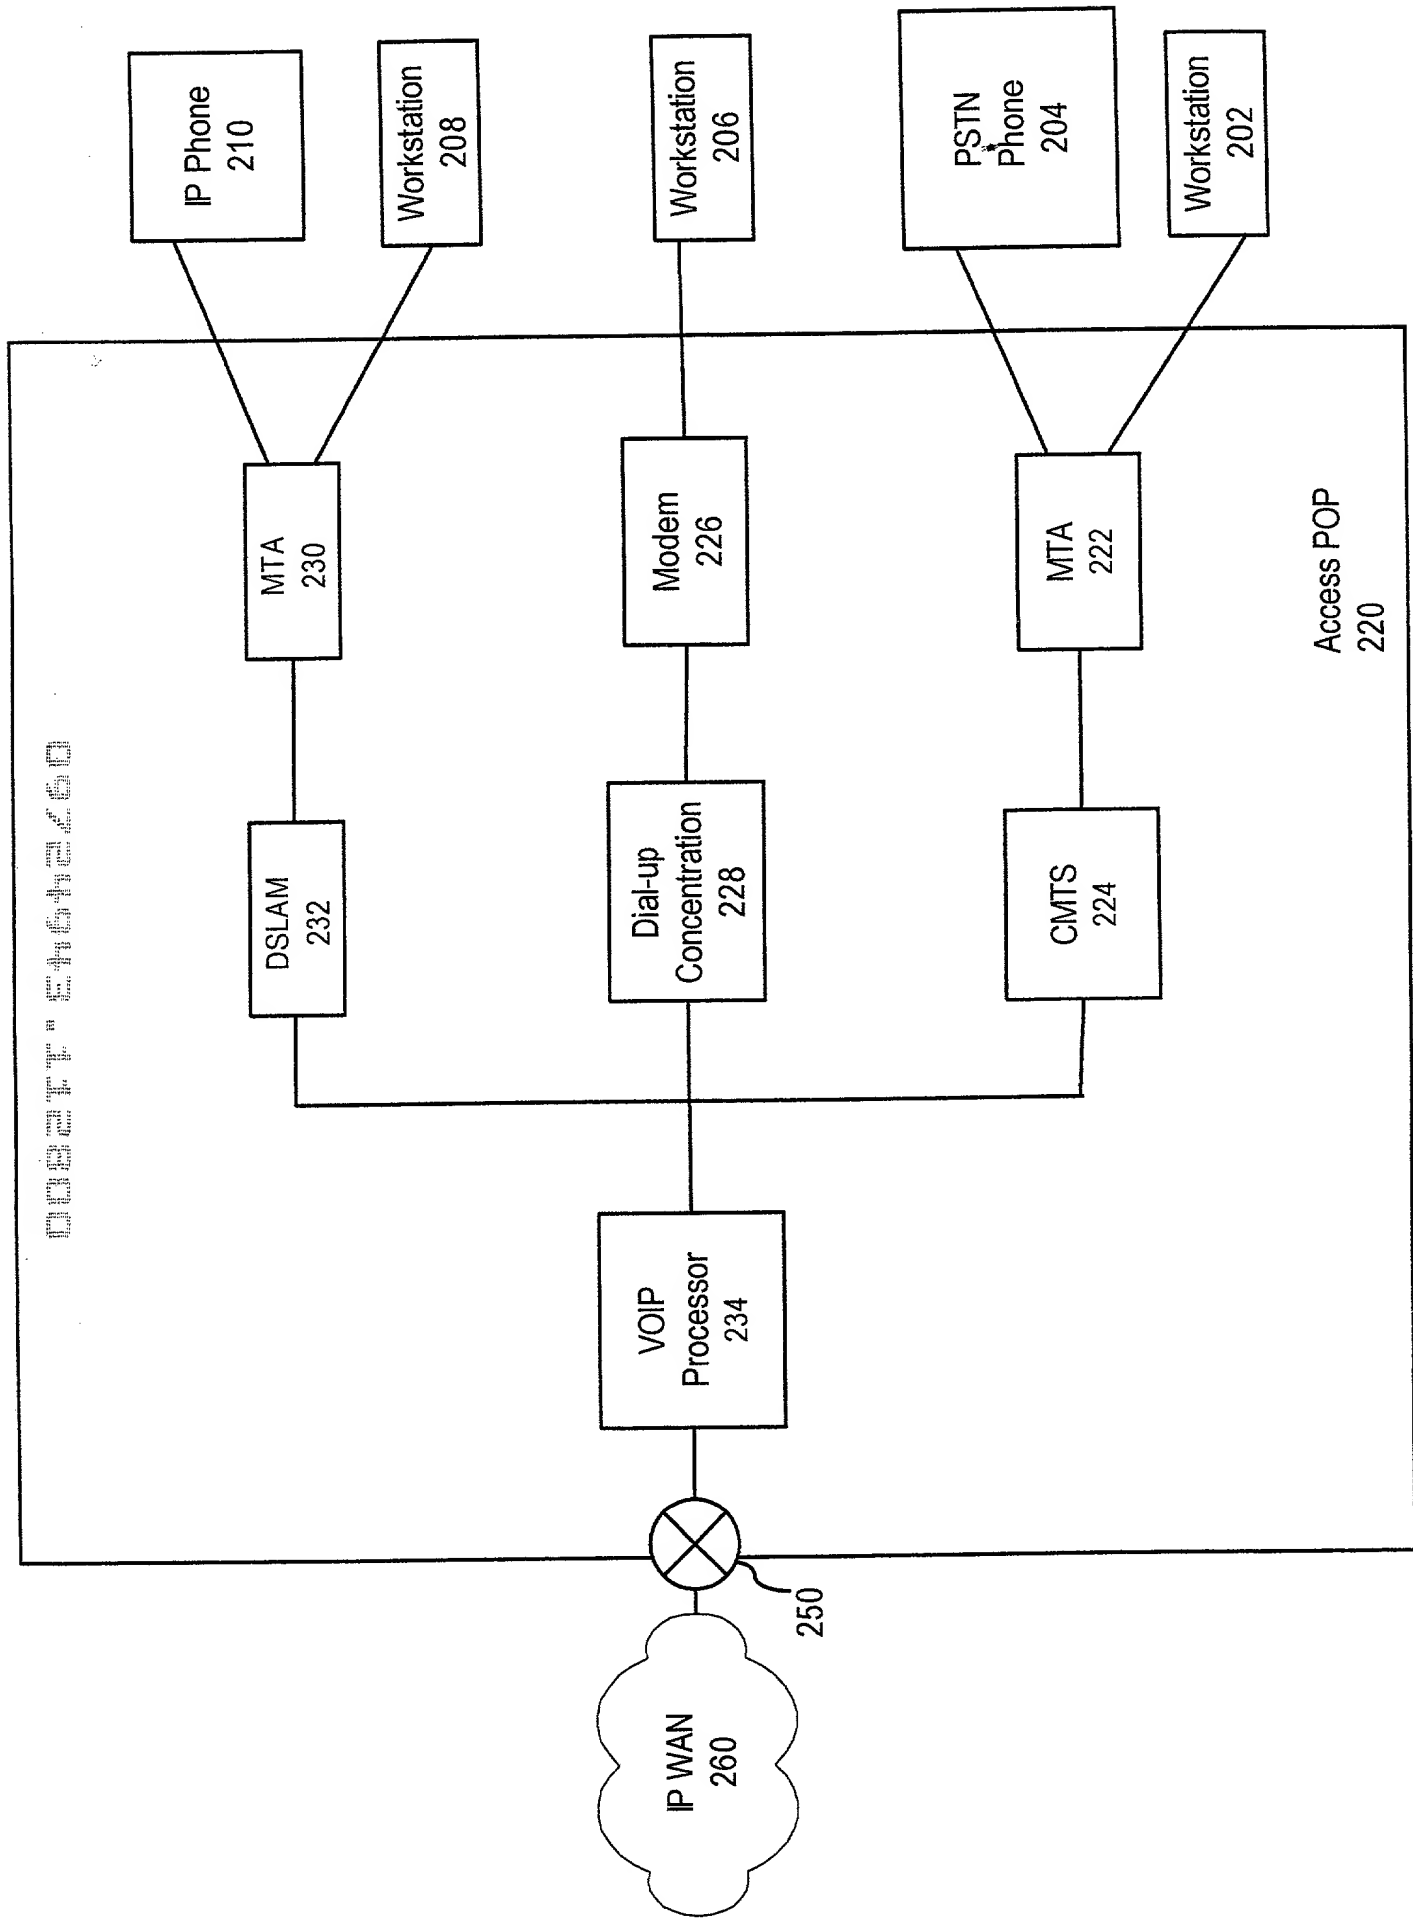
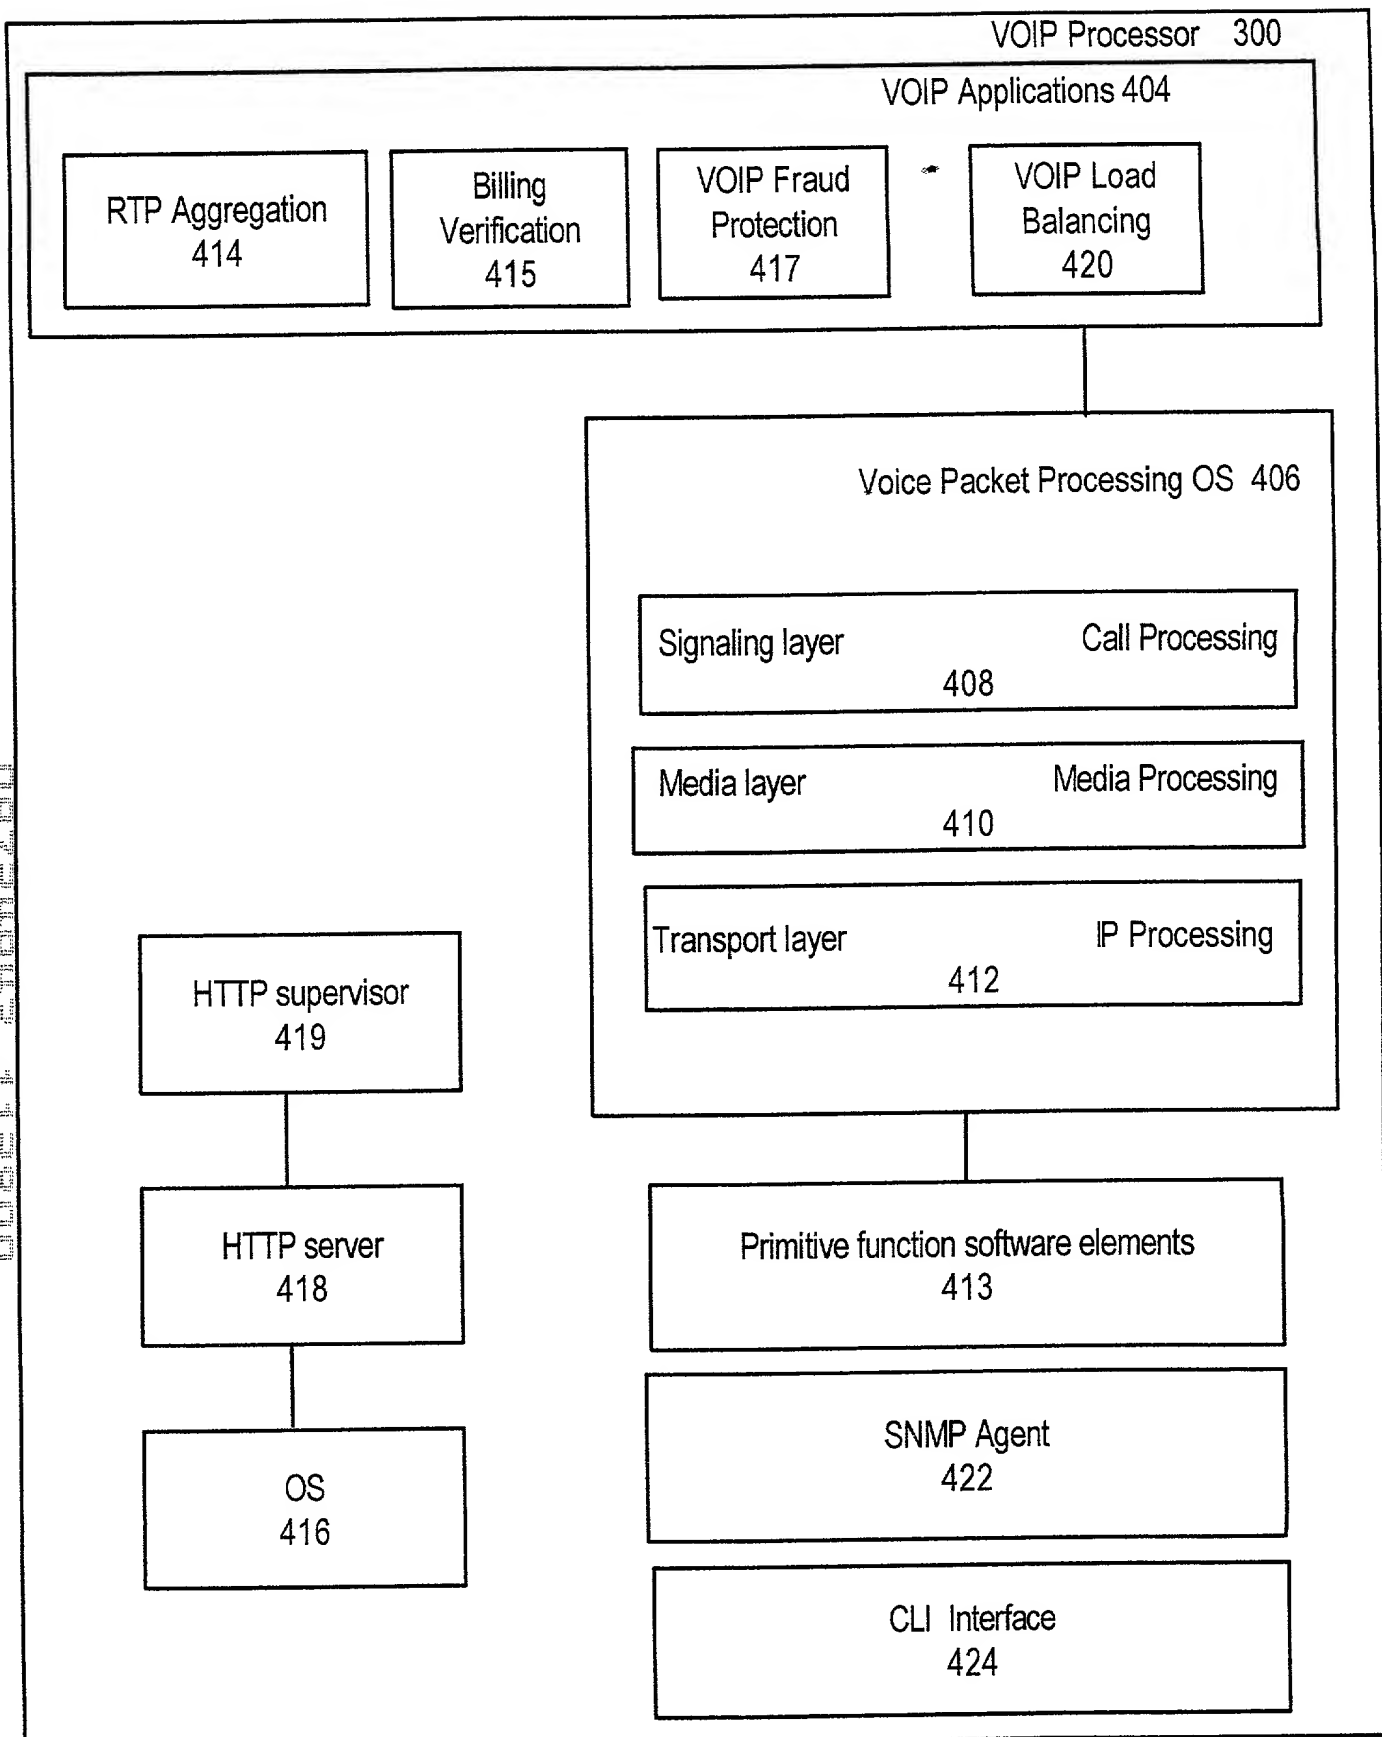


FIG. 2

The diagram illustrates the internal architecture of a network device 300. It features a large rectangular frame representing the device. On the left side, a large box is labeled "Switching Circuitry 302". On the right side, a vertical stack of components is shown: "Physical Interfaces 304" at the top, followed by "Classification Engines 306", then "Classification Tables 307", then "Network Processors 308", and finally "Host CPU 310" at the bottom. A vertical line connects "Classification Engines 306" to "Classification Tables 307". A curved line on the right side of the frame points to the label "300".

**FIG. 3**

000001 000002 000003 000004 000005 000006 000007 000008 000009 000010 000011 000012 000013 000014 000015 000016 000017 000018 000019 000020 000021 000022 000023 000024 000025 000026 000027 000028 000029 000030 000031 000032 000033 000034 000035 000036 000037 000038 000039 000040 000041 000042 000043 000044 000045 000046 000047 000048 000049 000050 000051 000052 000053 000054 000055 000056 000057 000058 000059 000060 000061 000062 000063 000064 000065 000066 000067 000068 000069 000070 000071 000072 000073 000074 000075 000076 000077 000078 000079 000080 000081 000082 000083 000084 000085 000086 000087 000088 000089 000090 000091 000092 000093 000094 000095 000096 000097 000098 000099 000100



**FIG. 4A**

## Primitive function software elements 414

### IP layer primitives 430

Packet dropping routine 432

Packet duplication routine 434

Packet marking routine 436

Label manipulation routine 438

Traffic policing routine 440

Packet scheduling routine 442

Packet re-routing routine 444

Tunneling routine 446

Encryption routine 448

Packet injection routine 450

Packet compression routine 452

### Media layer primitives 460

Transcoding routine 462

RTP aggregation routine 464

Header compression routine 466

Media modification routine 468

Media reconstruction routine 470

Media duplication routine 472

Media re-routing routine 474

### Signaling layer primitives 480

Protocol translation routine 482

Call detail record generation routine 484

Number translation routine 486

Drop call routine 488

Resource reservation routine 490

Call re-routing routine 492

Call re-direction routine 494

Call signaling duplication routine 496

Signaling aggregation routine 498

**FIG. 4B**

VOIP APPLICATIONS 404

RTP AGGREGATION APPLICATION 414

APPLICATION-LEVEL CALL ROUTING APPLICATION 431

SCALABLE TRAFFIC ENGINEERING APPLICATION 433

RSVP AGGREGATION APPLICATION 435

LOAD BALANCING APPLICATION 437

BILLING & INTERCONNECT VERIFICATION 439

ACCESS CONTROL/FRAUD DETECTION 441

RTP RECONSTRUCTION APPLICATION 443


CALL PRIVACY APPLICATION 445

DIFFERENTIATED SERVICES APPLICATION 447

LAWFUL INTERCEPTION APPLICATION 449

SERVICE SELECTION APPLICATION 451

**FIG. 4C**



**FIG. 5**

**FIG. 6**









FIG. 8



## DECLARATION AND POWER OF ATTORNEY

As below named inventors, we hereby declare that:

Our residence, post office and citizenship are as stated below next to our names,

We believe that we are the original, first and joint inventors of the subject matter claimed and for which a patent is sought on the invention METHOD AND APPARATUS FOR MONITORING AND PROCESSING VOICE OVER INTERNET PROTOCOL PACKETS, the specification of which

☒ is attached hereto.

☐ was filed on \_\_\_\_\_ as Application Serial No. \_\_\_\_\_ and was amended on (if applicable).

We hereby state that we have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

We acknowledge the duty to disclose information which is known to us to be material to patentability in accordance with Title 37, Code of Federal Regulations, Section 1.56.

We hereby claim foreign priority benefits under Title 35, United States Code, Section 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

**Prior Foreign Applications(s):**

Number	Country	Day/Month/Year filed	Priority Claimed
			<input type="checkbox"/>
			<input type="checkbox"/>

We hereby claim the benefit under 35 USC §119(e) of any United States provisional application(s) listed below.

**Prior Provisional Application(s):**

Application Number	Filing Date
60/226,207	August 18, 2000

We hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, We acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

**Prior U.S. Application(s):**

Serial No.	Filing Date	Status: Patented, Pending, Abandoned
------------	-------------	--------------------------------------

We hereby declare that all statements made herein of our own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

We hereby appoint the following attorney(s) and/or agent(s): Brian D. Hickman, Reg. No. 35,894; Christopher J. Palermo, Reg. No. 42,056; Bobby K. Truong, Reg. No. 37,499; Edward A. Becker, Reg. No. 37,777; Marcel K. Bingham, Reg. No. 42,327; Carl L. Brandt, Reg. No. 44,555; Carina M. Tan, Reg. No. 45,769; and Craig G. Holmes, Reg. No. 44,770, all of

HICKMAN PALERMO TRUONG & BECKER, LLP  
1600 Willow Street  
San Jose, California 95125-5106

with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith, and all future correspondence should be addressed to them.

Full name of sole or first inventor: Shai Mohaban

Inventor's signature:

Date:

Residence: 678 Durshire Way, Sunnyvale, CA 94087

Citizenship: Israel

Post Office Address:

Full name of second inventor: Itzhak Parnafes

Inventor's signature:

Date:

Residence: 20050 Rodrigues Avenue, Apt. B, Cupertino, CA 95014

Citizenship: Israel

Post Office Address:

Full name of third inventor: Opher Kahane

Inventor's signature:

Date:

Residence:

Citizenship:

Post Office Address: One Waters Park Drive, Suite 270, San Mateo, CA